Portfolio Management Under Epistemic Uncertainty Using Stochastic Dominance and Information-Gap Theory

D. Berleant[*], L. Andrieu[1], J.-P. Argaud[1], F. Barjon[1], M.-P. Cheong[2], M. Dancre[1], G. Sheble[2], and C.-C. Teoh[2]

[1]Electricité de France; [2]Portland State University; [*]principal contact, Dept. of Information Science, ETAS Building Rm. 258D, University of Arkansas at Little Rock, 2801 S. University Ave., Little Rock, Arkansas 72204, USA, berleant@gmail.com, phone (501) 683-7265, fax (501) 683-7049

Abstract

Portfolio management in finance is more than a mathematical problem of optimizing performance under risk constraints. A critical factor in practical portfolio problems is severe uncertainty – ignorance – due to model uncertainty. In this paper, we show how to find the best portfolios by adapting the standard risk-return criterion for portfolio selection to the case of severe uncertainty, such as might result from limited available data. This original approach is based on the combination of two commonly conflicting portfolio investment goals:

1) obtaining high expected portfolio return, and
2) controlling risk.

The two goals conflict if a portfolio has both higher expected return and higher risk than competing portfolio(s). They can also conflict if a reference curve characterizing a minimally tolerable portfolio is difficult to beat.

To find the best portfolio in this situation, we first generate a *set* of optimal portfolios. This set is populated according to a standard mean-risk approach. Then we search the set using stochastic dominance (SSD) and Information-Gap Theory to identify the preferred one. This approach permits analysis of the problem even under severe uncertainty, a situation that we address because it occurs often, yet needs new advances to solve. SSD is attracting attention in the portfolio analysis community because any rational, risk-averse investor will prefer portfolio y_1 to portfolio y_2 if y_1 has SSD over y_2. The player's utility function is not relevant to this preference as long as it is risk averse, which most investors are (e.g. De Giorgi 2005 [7], Berleant et al. 2005 [3]).

Keywords: Epistemic uncertainty, finance, imprecise probabilities, info-gap, information-gap, portfolio, probability boxes, stochastic dominance.

1 Introduction

A portfolio consists of a set of *segments*, each of which is predefined as a particular asset category, such as stocks, bonds, commodities, etc. Solving the selection problem means determining the best proportion each segment should be of the total investment. The portfolio selection problem is the subject of a vast body of work. The process can be divided into two phases. The first is asset allocation, in which investor philosophy, including risk position, is used to choose the best percentage of the portfolio to place in each segment. The second, rebalancing, responds to changes in asset values by adjusting the percentages so that the portfolio continues to accurately reflect the investment philosophy. This work focuses on allocation. The well known CAPM leads to various allocation strategies, including for example BIRR and BARRA

(search the Web for further information about these). The correct treatment of the risk-reward problem addressed by Markowitz (1952 [12]) is fundamental to such modern methods, and its extension to problems characterized by severe uncertainty motivates this report.

Little has been done to determine portfolio allocation when dependency relationships, such as correlations, among portfolio segment return distributions are unknown. We address this problem with a novel application of Information-Gap Theory (Ben-Haim 2006 [2]), using it together with the concept of *second-order stochastic dominance* (SSD) to help choose among portfolio allocations. SSD holds between two distributions r_1 and r_2 when the curves of their integrals do not cross. The slower rising curve is then said to have second-order stochastic dominance over the other curve. If r_2 has SSD over r_1, then we write $r_2 \succ_2 r_1$. Analogously, FSD (first-order stochastic dominance, $\prec_1$) applies if the distribution curves themselves do not cross. However, most investors are risk averse, and if $r_2 \succ_2 r_1$ then any risk averse player will prefer r_2 (e.g. Perny et al. 2007 [13]). FSD is thus an unnecessarily strong (and therefore undesirable) constraint for the risk averse player.

We build on a standard approach to finding optimal portfolios based on *mean* and *risk* and parameterized by amount of risk aversion, arising originally from Markowitz (1952 [12]). The mean is the expectation (i.e. average) for a distribution describing the investment return, while the risk expresses the danger of loss or low returns and is typically a measure of the spread of the return distribution. Optimal portfolios are identified by finding the *weights* of the portfolio segments such that a *mean-risk* objective function is maximized (e.g. Elton et al. 2003 [8]). "Mean-risk" refers to a tradeoff between seeking a return distribution with a high mean, which is good, while minimizing the higher risk that tends to be associated with a high mean return, which is bad.

Formally, the problem to be considered is to find such a portfolio given the constraints

$$r = \sum_{i=1}^{s} w_i r_i \succ_2 \tilde{R} \qquad (1)$$

$$\sum_{i=1}^{s} w_i = 1 \qquad (2)$$

where r is a portfolio return distribution, i is one of the s segments in the portfolio, w_i is the weight of segment i, and r_i is the return distribution of segment i. $\tilde{R}$ is a given reference curve representing a minimally tolerable bound (the "risk limit") that the return distribution should not cross. As an additional constraint set (Eq. 2), segment weights sum to 1 because each weight is a proportion of the whole. Each weight may be required to be within some interval in order to enforce a balance across segments, as might be specified by a company's business model constraints and investment policies.

In current practice optimization would typically be done without the $\succ_2$ constraint, but perhaps with other constraints present such as Value-at-Risk (VaR), which is known to be mathematically inconsistent, or Conditional VaR (CVAR) which is less intuitive but without VaR's consistency problem [3].

A given portfolio's return distribution can be tested for compliance with an SSD constraint using numerical integration. Numerical integration can be done straightforwardly by summing areas of trapezoids under the curve. The size and number of trapezoids to sum is determined by the step

size chosen for the integration process. Given a step size, SSD is considered to hold if the summed areas of all trapezoids to the left of any given value x_i are lower for the return distribution r of a candidate portfolio than for reference curve $\tilde{R}$. A set of representative return values x_1, x_2, x_3, …, x_n that are possible sample value of r should be checked. These points should cover low and high values of return, as well as a reasonable number of intermediate points (e.g. m = 10 or 20).

An optimal portfolio might or might not comply with an additional requirement that it have stochastic dominance over a reference curve $\tilde{R}$. A second-order stochastic dominance relation (" $\succ_2$ ") between two distributions ensures that the dominant portfolio is preferred by any risk averse player (De Giorgi 2005 [7]). Risk aversion implies that, given two return distributions with the same expected return, the one with less spread (e.g. variance) is preferred. Define *robustness* as the amount by which a portfolio dominates by SSD a reference curve (robustness could be negative if it does not dominate). By testing various optimal portfolios for robustness, one with the highest robustness available can be identified. Alternatively, one with the highest expected return that also meets the SSD constraint could be found. In either case, the strategy is to search among a set of optimal portfolios provided, even if barely, by an under-constrained optimization problem for the one that is best according to a second criterion. This is discussed next.

2 Searching the optimal portfolios for the best one

The first step is to generate a set of optimal portfolios to search within for the best. In the standard approach of Markowitz portfolio theory, the desirability of a portfolio return distribution r given a value of a parameter z describing the degree of aversion to risk is:

$$f(z,r) = mean(r) - z * risk(r). \qquad (3)$$

We can build on the concept of Eq. 3 by defining a function identifying the desirability of the best portfolio(s) for a given risk position z:

$$OPT(z) = \sup_{y \in Y}(\mu_y - z\sigma_y^2), z \geq 0 \qquad (4)$$

where $z \geq 0$ means that the investor is not risk-seeking, Y is the set of portfolios y complying with constraints (1) and (2), and μ_y and σ_y^2 are the expected return and variance respectively of the return of portfolio y. The gist of (4) may be used to define the set of optimal portfolios as:

$$OPT = \{ y : (y \in Y) \wedge (\forall y' \in Y : \exists z > 0 : \mu_y - z\sigma_y^2 \geq \mu_{y'} - z\sigma_{y'}^2 \} \qquad (5)$$

where $z > 0$ restricts the set to portfolios that are optimal for some risk-averse party. This restriction is consistent with most investors, and also with our focus on SSD, which is only valid for ordering the desirability of portfolios if the party is risk averse. It is possible to account for properties of portfolio variation besides variance (Chabi-Yo 2004 [5]), but σ^2 is nevertheless widely used to model risk position. The variance approach can be elaborated to model different variances for different portfolio segments, and for the covariances between segments. The variance of the return distribution for a portfolio y then follows $\sigma_y^2 = \sum_{i=1}^{s}\sum_{j=1}^{s} w_i w_j \sigma_{ij}$, where w_i and w_j are the weights of portfolio segments i and j, and σ_{ij} is their covariance (when $i=j$, $\sigma_{ij} = \sigma_{ii} = \sigma_i^2$). Even in this more sophisticated model, z is still a simple scalar coefficient.

Thus to generate a set of optimal portfolios, $OPT(z)$ can be evaluated repeatedly with different values of z (Figure 1).
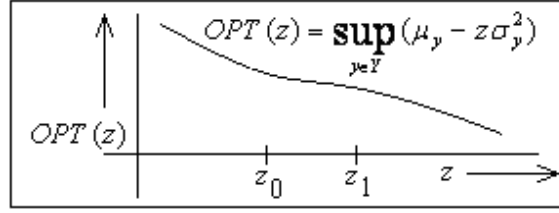

Figure 1. Each value of z implies some maximum value of $OPT(z)$ resulting from some portfolio(s) in the feasible set.

2.1 $OPT(z)$ and the efficient frontier

Let us now relate Figure 1 (representing an optimal tradeoff between mean return and risk) to the concept, due to Markowitz, of the *efficient frontier* (Figure 2). The efficient frontier is plotted as a curve on a plane with axes for mean and variance. Each point on the curve represents a portfolio that is optimal in that no other portfolio has a return distribution with the same (or higher) mean and lower variance, or the same (or lower) variance and higher mean. Formally, a portfolio y is on the efficient frontier if

$$\forall y' \in Y, \neg \exists y': (\mu_{y'} \geq \mu_y \wedge \sigma_{y'}^2 < \sigma_y^2) \vee (\mu_{y'} > \mu_y \wedge \sigma_{y'}^2 \leq \sigma_y^2).$$
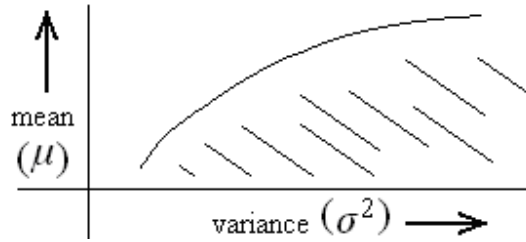

Figure 2. An example of part of an efficient frontier.

Because of the well-known virtues of the efficient frontier in portfolio analysis [8, 12], it is useful to establish that points on the $OPT(z)$ curve (Figure 1) are also on the efficient frontier (Figure 2), and that points on the efficient frontier are also on the $OPT(z)$ curve.

Let μ_i, $i = 1, \ldots, s$ be the mean return of portfolio segment i. Let different segments be characterized by corresponding random variables r_i. For every set of s real valued weights w_i, with each $w_i \in [0, 1]$, we have $\sigma_y^2 = \sum_{i=1}^{s} \sum_{j=1}^{s} w_i w_j \sigma_{ij} = E[(r_y - E[r_y])^2]$, where $r_y = \sum_{i=1}^{s} w_i r_i$ is the return distribution of the corresponding portfolio y. Therefore, we always have $\sum_{i=1}^{s} \sum_{j=1}^{s} w_i w_j \sigma_{ij} \geq 0$, i.e., the matrix of σ_{ij} is semi-definite.

Theorem.
- If a portfolio y is on the efficient frontier then it solves $OPT(z)$ for some z (see Eq. 4).
- If a portfolio y belongs to set OPT, it is on the efficient frontier (see Eq. 5).

4

We start with a proof of the second statement by contradiction. Assume that portfolio $y \in OPT$ is not on the efficient frontier. Then, $\exists y': (\mu_{y'} \geq \mu_y \wedge \sigma_{y'} < \sigma_y) \vee (\mu_{y'} > \mu_y \wedge \sigma_{y'} \leq \sigma_y)$. In that case, $\mu_{y'} - z\sigma_{y'}^2 > \mu_y - z\sigma_y^2$, contradicting the definition of OPT (see Eq. 5).

Now consider the first statement. Assume point (μ_y, σ_y^2), representing portfolio y, is on the efficient frontier (Figure 2). Consider the set P of points (μ, σ^2) for which $\mu \leq \mu_y \wedge \sigma^2 \geq \sigma_y^2$ for some portfolio y. We say these points are *p-dominated* by y. The matrix of the second derivatives of the function σ_y^2 is positive definite, so the function is convex. If we have two portfolios y' and y'', then for their convex combination $y(w) = wy' + (1-w)y''$ we have $\mu_{y(w)} = w\mu_{y'} + (1-w)\mu_{y''}$, but $\sigma_{y(w)}^2 \leq w\sigma_{y'}^2 + (1-w)\sigma_{y''}^2$ (note the $\leq$ sign) because the variance of a sum may be less than the sum of the variances. Thus the linear combination of $(\mu_{y'}, \sigma_{y'}^2)$ and $(\mu_{y''}, \sigma_{y''}^2)$, that is, $(w\mu_{y'} + (1-w)\mu_{y''}, w\sigma_{y'}^2 + (1-w)\sigma_{y''}^2)$, is p-dominated by portfolio $y(w)$ and thus is in P.

In general, if points $p_1 = (\mu_1, \sigma_1^2)$ and $p_2 = (\mu_2, \sigma_2^2)$ are in P, there exist portfolios y' and y'' which p-dominate them respectively. Therefore $wp_1 + (1-w)p_2$ is p-dominated by $(w\mu_{y'} + (1-w)\mu_{y''}, w\sigma_{y'}^2 + (1-w)\sigma_{y''}^2)$, which we already showed is p-dominated by portfolio $wy' + (1-w)y'' = y(w)$. By transitivity we conclude that $p = wp_1 + (1-w)p_2$ is p-dominated by $y(w)$, and hence is in P.

Thus set P contains a convex combination of every two of its points and therefore is convex. We can also prove that P is closed. Indeed, the set of all the portfolios is bounded and closed hence compact. If $p_n = (\mu_n, \sigma_n^2)$ is p-dominated by some portfolio y_n and $p_n \to p'$ then from the sequence y_n we can extract a convergent subsequence $y_{n_k} \to y$. Because p_{n_k} is p-dominated by y_{n_k} we can conclude that, as $k \to \infty$, p' is p-dominated by y. Therefore $p' \in P$.

Portfolio y, earlier assumed to be on the efficient frontier, is clearly in P. It cannot be in the interior of P because if it was, a sufficiently small neighborhood around it would also be in P, and this neighborhood would contain points with higher μ and lower σ^2. Since every point in P is p-dominated by some portfolio, this portfolio would therefore have higher μ and lower σ^2 than y, contradicting our assumption that y is on the efficient frontier.

Therefore y is on the boundary of a closed convex set P. It is known (e.g. Rockafeller 1970 [14]) that in this case, there exists a line $a\mu_y - b\sigma_y^2 = c_a$ which goes through y such that all points in P are on one side of the line. Dividing by a (for $a \neq 0$) we get the line in the desired form $\mu_y - z\sigma_y^2 = c$ (the case $a = 0$ corresponds to $z = \infty$). Portfolio y is on this line, while all the others are on the other side. We thus conclude that y indeed is on curve $OPT(z)$.

Since the curve of Figure 1 contains the points representing the portfolios we need to consider, we can proceed to analyze the curve of Figure 1 to find the best point. We illustrate this with an example next.

2.2 Example

Consider a 3-segment portfolio problem in which we seek to optimize the weights of the segments. The return distribution for segment one is normal: $r_1 \sim Normal(1.1, 0.25)$. The distribution for segment two is exponential: $r_2 \sim Exponent(1.0, 1.0)$. Finally the distribution for segment three is uniform: $r_3 \sim Uniform(1.2, 0.48)$. Corresponding interval restrictions on the weights of the segments are $w_1 \in [0.2, 0.3]$, $w_2 \in [0.4, 0.6]$ and $w_3 \in [0.2, 0.3]$.

Taking into account a full range of assets, the degree of risk aversion, z, for many investors ranges from 2 to 4 (Grossman and Shiller 1981 [10]). A default value of $z = 3$, representing a typical degree of risk aversion (Bodie et al. 1999 [4]), implies that a relatively cautious investor will have z > 3, while a relatively aggressive investor will have z < 3. In practice, such typical textbook values might not entirely cover the full range of degrees of risk aversion we would like to deal with, so we examined z values for this example over the broader range of $z = 0.2$ to $z = 5$. The optimal portfolios found for various values of z are shown in Table 1. The return distributions for the two portfolios that were optimal for extremes $z = 0.2$ and $z = 5$ are shown in Figure 3. (To reduce clutter, curves for intermediate values of z are not shown).
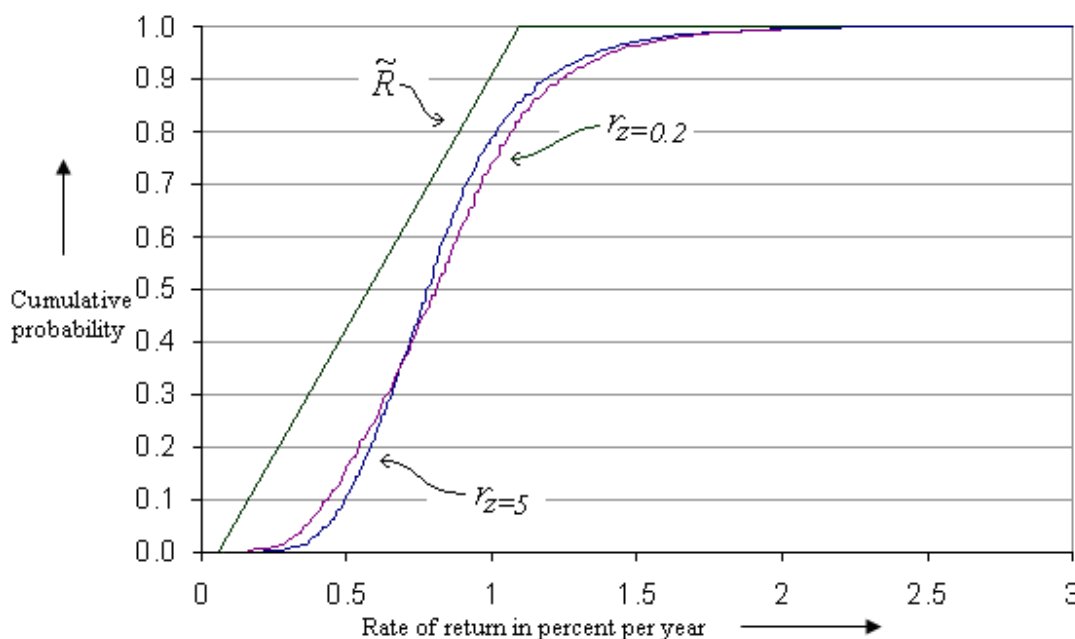
End of example

Figure 3. Distributions of the returns of the two portfolios that are optimal for $z = 0.2$ (shallower curve) and $z = 5$. Both have SSD over the sample reference curve shown.

2.2.1 What's next. Finding the best portfolio from the optimal portfolios on the efficient frontier is an investment decision. While it would be a simple one to make by specifying a value for z, most investors would not be able to state (or accept) any specific z value as a firm constraint on their portfolios, so it is not a common decision criterion. Thus investigation of other criteria for comparing portfolios is important in reaching conclusions about portfolio composition. Some potential criteria are relatively straightforward. Others might be better but not as simple. Identifying what criteria for portfolio quality are best under what conditions presents an interesting problem. In this report we discuss and compare some criteria and related issues based on SSD and Info-Gap Theory.

Value of z	Optimal portfolio segment weights		
	s_1	s_2	s_3
0.2	0.2	0.5	0.3
0.4	0.2	0.5	0.3
0.6	0.2	0.5	0.3
0.8	0.2	0.5	0.3
1	0.2	0.5	0.3
1.2	0.2146	0.5	0.2854
1.4	0.2316	0.5	0.2654
1.6	0.2432	0.5	0.2568
1.8	0.2527	0.5	0.2473
2	0.2603	0.5	0.2397
2.2	0.2665	0.5	0.2335
2.4	0.2717	0.5	0.2283
2.6	0.2761	0.5	0.2239
2.8	0.2798	0.5	0.2202
3	0.2831	0.5	0.2169
3.2	0.286	0.5	0.214
3.4	0.2885	0.5	0.2115
3.6	0.2907	0.5	0.2093
3.8	0.2927	0.5	0.2073
4	0.2945	0.5	0.2055
4.2	0.2959	0.5	0.2041
4.4	0.2976	0.5	0.2024
4.6	0.299	0.5	0.201
4.8	0.3	0.5	0.2
5	0.3	0.5	0.2

Table 1. Optimal portfolios (i.e. optimal segment weight vectors) for various values of z. Return distributions of the segments are assumed to be independent.

2.3 Criteria for choosing the best portfolio

Each criterion named in items **1** through **4c** of Table 2 is described and discussed in its own subsection. Basic notation is explained next, while technical details about the meaning of the horizontal dimension of Table 2 ("***Quality Metric***") and its vertical dimension ("***Objective***") depend on the cell of the table under consideration and thus are given later in the detailed item descriptions.

Notation:
- μ is an expected (i.e., mean) return.
- |SSD| is the *amount* by which one distribution dominates another, using as a criterion second-order stochastic dominance. This is the minimum horizontal distance between the integrals of two cumulative distributions (Figure 4). In other words, |SSD| measures how much one curve for the integral of a distribution can be moved toward another one along the x axis before the two curves touch. |SSD| formalizes the amount of separation between the integrals of two distributions.

- α is a parameter expressing the amount of ignorance (sometimes called epistemic uncertainty) about the shape of a return distribution. Specifying $\alpha = 0$ will lead us to use the best-guess estimate of its shape, while $\alpha = 1$ will lead to a distribution that incorporates pessimistic assumptions about possible errors in the shape of the best-guess distribution. Intermediate values of α are then used to generate distributions by interpolating between the distributions implied by $\alpha = 0$ and $\alpha = 1$ using "horizontal averaging" (defined later when it is used). Our use of α is an example of the more general definition of α in Info-Gap Theory (Ben-Haim 2006 [2]).
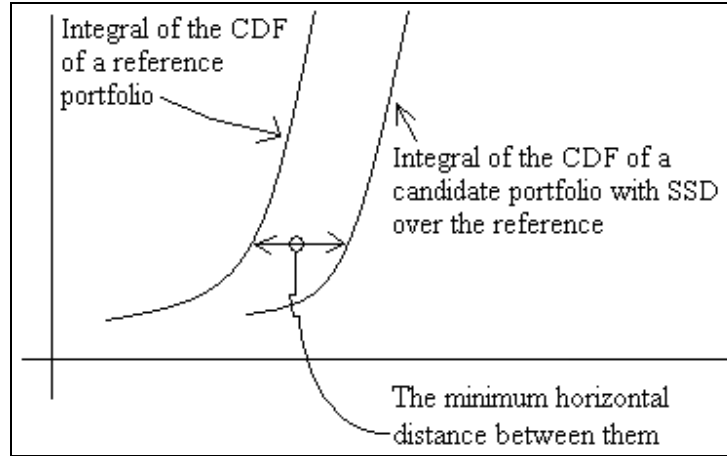

Figure 4. Two curves for the integrals of distributions, and the separation between them. Recall that the integral of a density function (PDF) is a cumulative distribution (CDF). The integral of the CDF is what is used in determining SSD.

		Quality Metric	
		|SSD|	α **(alpha)**
O b j e c t i v e	**Maximize Robustness (to achieve secure performance)**	**1.** Given $\tilde{R}$, find the portfolio(s) with the highest |SSD| over $\tilde{R}$, i.e., move $\tilde{R}$ to the right until further movement would disqualify every portfolio.	**3.** Given $\tilde{R}$, find a portfolio with the highest possible α.
	Maximize μ (to achieve best performance within the risk limit)	**2.** Find a portfolio y with return distribution r_y such that $r_y \succ_2 \tilde{R}$, choosing one with the highest possible mean return μ.	**4a.** Find a portfolio with the highest mean return μ from among those with SSD over reference curve $\tilde{R}$ for *any* dependency relationships among segments.
			4b. Generalize 4a by requiring SSD for only *some* dependencies. The precise meaning of "some" is determined by the value of α.
			4c. Find the demand value of information about α in order to choose what value of α to use in 4b.

Table 2. Approaches to finding a best portfolio. $\tilde{R}$ is a reference curve.

2.3.1 Computational complexity and searching the set of candidate portfolios. All approaches enumerated in Table 2 involve searching the set of candidate portfolios corresponding to points on the curve of Figure 1. To initially populate the set with *n* members requires *n* corresponding evaluations of *OPT*(*z*), so the time required to populate is *n* × the time required to evaluate *OPT*(*z*). The value of *n* is the number of points on the curve of Figure 1 for which corresponding portfolios are to be computed. To keep *n* low enough to make the problem tractable, we used a heuristic approach to searching the curve for a point corresponding to a best portfolio. The approach will find the maximum if the search space is sufficiently well behaved, otherwise it might get stuck at a local maximum. There is no indication that ill-behaved search spaces would be common, but if such a space was present a more comprehensive search strategy would be needed to find the maximum.

Algorithm for searching for the optimal portfolio.
1. Given a measure of portfolio quality, evaluate portfolios corresponding to *n* points on the curve of Figure 1, selected to be spaced at representative values of *z* from the minimum to the maximum value of interest. In this work we chose 6 values for *z* of 0.2, 1, 2, 3, 4, and 5, but the precise value of *n* is unimportant.
2. Let z_k be the value of *z* for the best portfolio of those tested. Next test the portfolio that is optimal for $z = z_k + \varepsilon$, for ε small but not so small that error due to the limited precision and accuracy of machine arithmetic is significant. If the optimal portfolio for $z = z_k + \varepsilon$ is better than the optimal portfolio for $z = z_k$, then search for even better portfolios in the range $[z_k, z_{k+1}]$. Otherwise, search for better portfolios in the range $[z_{k-1}, z_k]$.
3. Define *n* equally spaced values within the new range for *z* that we wish to search. Starting from the value nearest to z_k, test optimal portfolios corresponding to values of *z* progressively further from z_k, stopping at the value for *z* within the range that corresponds to the best portfolio.
4. Loop back to Step 2, or stop if significant further improvement seems unlikely.

For *n* values of *z* per iteration, *i* iterations, time *t* to evaluate *OPT*(*z*) and test the quality of the portfolio it returns, the worst case run time for the algorithm is $n \cdot i \cdot t$.

In the following sections we discuss each criterion in Table 2, starting from item **1** and proceeding through **4c**.

2.3.2 Given $\tilde{R}$, find the portfolio(s) with the highest |SSD| over $\tilde{R}$ (Table 2, item 1). Given a minimum acceptable return distribution $\tilde{R}$ (the "reference" curve), an investor may wish to find a portfolio that is as sure as possible to be better, despite errors in its estimated return distribution. For a risk-averse investor, if $y_r \succ_2 \tilde{R}$ then y_r is better (Section 1). For such an investor a natural goal is to find the portfolio whose return has the highest obtainable |SSD| over $\tilde{R}$ (Figure 4).

While this method explicitly seeks to optimize based on SSD (and not expected return), the resulting portfolio will often also have a relatively high expected return, simply because maximizing |SSD| tends to favor curves that are to the right of curves with less |SSD|, and hence have higher means. Although this tendency would be welcomed by an investor, the point of this method is actually to maximize robustness in the sense of assuring that the distribution of the chosen portfolio really does stochastically dominate the reference curve, even if the distribution is inaccurately stated. Such inaccuracy is a genuine concern because of the incompleteness of information available about future returns.

Example (continued). Building on the example given previously (Section 2.2), portfolios corresponding to points on the curve of Figure 1 for representative values of z were computed. For each, the return distribution was compared to a reference curve. Results are shown in Table 3. For the reference curve used, the best portfolio of those tested was that corresponding to $z = 4$. The composition of that portfolio was shown in Table 1.

z	|SSD|	μ
0.2	0.2662	1.0800
1	0.2662	1.0800
2	0.2841	1.0740
3	0.2872	1.0717
4	**0.2886**	1.0705
5	0.2864	1.0700

Table 3. For the given $\tilde{R}$, \|SSD\| is highest for $z = 4$ than for $z = 0.2, 1, 2, 3,$ or 5. (However the expected return μ for $z = 4$ is not highest.)

2.3.3 Of the portfolios with SSD over $\tilde{R}$, find one with the highest possible μ (Table 2, item 2).

The objective here is to maximize *mean return* (μ), instead of \|SSD\| as was done in the previous section. This goal makes sense if an investor is more concerned about return than about possible errors in estimating return distributions. The idea is to search among portfolios with SSD over reference curve $\tilde{R}$ to find one with the highest expected return. The obvious benefit to choosing maximum expected return over maximum degree of SSD as an objective is its financial advantage. On the other hand, the rejected portfolio might actually be preferable due to its greater margin of safety, in the sense that its dominance is more robust to inaccuracies in the portfolio return distributions.

Maximizing μ suffers from another potential pitfall as well. An analyst might succumb to the temptation to try improving results by replacing reference curve $\tilde{R}$ with an $\tilde{R}'$ such that $\tilde{R}' \succ \tilde{R}$, because $\tilde{R}'$ is tougher to beat. This will tend to shrink the set of qualifying portfolios since it is likely that fewer will stochastically dominate $\tilde{R}'$ than $\tilde{R}$. This in turn will tend to reduce the highest μ available within that set: $\{y : r_y \succ \tilde{R}'\} \subseteq \{y : r_y \succ \tilde{R}\} \rightarrow \sup_{r_y \succ \tilde{R}'} \mu_y \leq \sup_{r_y \succ \tilde{R}} \mu_y$, where y is a possible portfolio, and r_y and μ_r are the return distribution and expected return, respectively, for portfolio y. In Figure 5, that means disqualifying portfolio y_2 because it fails to dominate $\tilde{R}'$ despite having a higher expected return than y_1, which does dominate $\tilde{R}'$. (FSD is illustrated for graphical clarity, but if the curves were all integrated then the analogous situation for SSD would be depicted.)

Example. Continuing the example developed earlier, optimal portfolios for selected values of z were analyzed to determine their mean returns μ in addition to their \|SSD\|. Whereas in Section 2.3.2 the portfolio corresponding to $z = 4$ was chosen because it has the highest \|SSD\| (Table 3, middle column), here we would choose the portfolio corresponding to $z = 0.2$ and $z = 1$ because, of those portfolios with \|SSD\| > 0, it has the highest expected return μ (Table 3, right column). The composition of that portfolio was shown in Table 1.
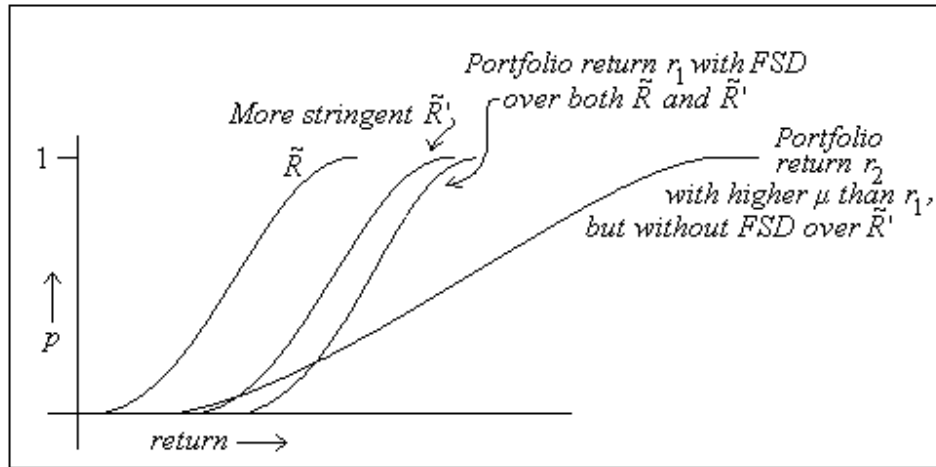
Figure 5. Shifting reference curve $\widetilde{R}$ to the right, yielding $\widetilde{R}'$, presents a more difficult FSD (first order stochastic dominance) requirement for a candidate portfolio to meet.

2.3.4 Given $\widetilde{R}$ find the portfolio(s) with the highest α (Table 2, item 3). We introduce this method by comparing it with a method described earlier. We then formalize it, and close with an example.

Introduction. Let us compare this approach to that of finding the portfolio with the highest |SSD| over a reference curve $\widetilde{R}$ (Table 2, item **1**; Section 2.3.2). Maximizing |SSD| is suitable when the goal is to make the chosen portfolio robust, in that its return distribution r will tend to have SSD over reference curve $\widetilde{R}$ even if its true distribution is somewhat different from the best estimate we have for r. A similar goal is supported by using the Information-Gap Theory uncertainty parameter, α, instead of |SSD|. However although α might be less immediately intuitive than the amount of second order stochastic dominance |SSD|, α can be better for modeling how much a distribution stochastically dominates a reference curve.

A significant problem with |SSD| as a measure of robustness, that α can circumvent, is that |SSD| measures robustness by determining how much the integral of a distribution can be shifted on the x-axis without intersecting a reference curve (Fig. 4). This metric is problematic. The accuracy of a return distribution depends on understanding of the influence of various factors such as leverage, reliability of historical data, expert judgment, and unspecified dependency relationships among asset prices. In general, eliminating inaccuracies due to such factors would not necessarily update a return distribution from $r_1 = F(x)$ to $r_2 = F(x+k)$, i.e., shift the distribution on the x-axis without changing its shape. But this is exactly what |SSD| assumes. Further, since estimates about the dependency relationships among the portfolio segments can be a major source of inaccuracy, it is desirable to account specifically for this source of error. These needs can be met using the α parameter to combine the concept of robustness (Ben-Haim 2006 [2]; Cheong et al. 2004 [6]) with that of bounding the family of distributions corresponding to the space of different possible dependency relationships among the return distributions of portfolio segments (e.g. Zhang and Berleant 2005 [15]).

To combine these concepts, consider first an example, the distribution for portfolio return labeled "**Best-guess return**" in Figure 6. This distribution is associated with some portfolio with segments, segment weights and a particular set of dependencies among the segments (perhaps that they are independent). Next, suppose one drew a separate distribution curve for each

conceivable dependency relationship among the segments. The resulting set of curves could be bounded by "envelope" curves, labeled "**Left envelope, dependencies unknown**" and "**Right envelope, dependencies unknown**" in Figure 6. Such envelopes can be obtained using Statool (Zhang and Berleant 2005 [15]) or RiskCalc (Ferson 2002 [9]); others have built ad hoc calculating software (Helton and Oberkampf 2004 [11]). Expressing return using these envelopes represents a refusal to make assumptions about the dependency relationships among the different portfolio segments. Such a strategy would make sense when dependencies are, in fact, unknown, as might be the case if little historical data exists.
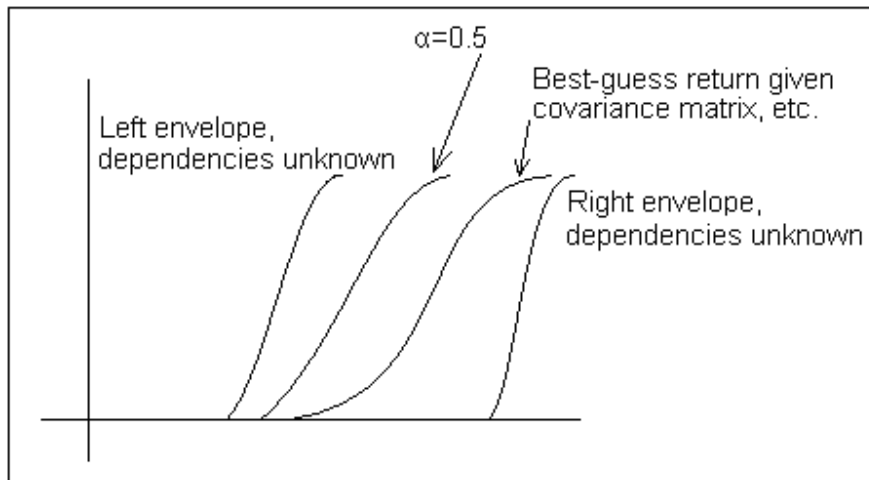

Figure 6. The "**Best-guess return**" distribution for the return of a given portfolio is the return distribution assuming all information used to construct it is correct.

Following the conventions of Info-Gap Theory (Ben-Haim 2006 [2]) we quantify the amount of epistemic uncertainty ("ignorance") with a parameter called α. Info-Gap Theory can use the value of α to generate error bounds in model outputs. These bounds define a space of possible results. If the worst case member of that space is an acceptable result, the model is robust against the amount of ignorance expressed by that value of α. If the model is not robust, then efforts to reduce ignorance could make it robust, because lower ignorance would be expressed with a smaller value of α, implying narrower error bounds. The consequent smaller space of possible results might be robust where the larger space was not.

Let the condition of zero epistemic uncertainty (i.e., $\alpha = 0$) correspond to full specification of the return distribution of a portfolio, such as the "**Best-guess return**" depicted in Figure 6. This curve might, for example, be provided by a financial analyst. It could be the distribution corresponding to a default assumption that the portfolio segments are independent. Alternatively, it might assume a specific covariance matrix stating dependencies among the portfolio segments derived by mining data on previous performance of the segments.

As suggested by its "Best guess" designation, the default distribution is not necessarily the actual distribution. Estimating portfolio segment distributions, and therefore the return distribution of an overall portfolio, may be done using historical data, economic projections, expert judgement, etc., but conclusions will always be error-prone because information is almost always limited and the future almost always contains a major element of unpredictability. Thus it is desirable to specify the uncertainty associated with its shape. Earlier in this section we assigned $\alpha = 0$ to designate a distribution for return with zero uncertainty about its shape: the "**Best-guess return**" curve. Now, let us assign $\alpha = 1$ to express some desired maximum amount of uncertainty about

the shape of the curve. Values for α between 0 and 1 express intermediate amounts of epistemic uncertainty.

The details of what bounding envelopes around the best-guess curve correspond to a given value of α depend on the details of the epistemic uncertainty one wishes to model. If epistemic uncertainty is due to ignorance about the dependency relationships among the return distributions of the portfolio segments, a possible modeling strategy would be,

1)   seek envelopes bounding the space containing all return curves corresponding to any mathematically possible dependency relationship among the segment return distributions (corresponding to $\alpha = 1$), and then,

2)   interpolate between the best guess curve and the envelopes to obtain more constraining bounds, which we will term "sub-envelopes," nested within the envelopes, which bound families of distributions corresponding to values of α for $0 < \alpha < 1$.

Ultimately, we must be able to determine (*a*) the maximum α for which a particular portfolio return distribution is sure to have stochastic dominance over a reference curve $\tilde{R}$, then (*b*) which of a set of candidate portfolios has a distribution with the highest maximum α. These ideas are described more formally next.

Formalization. The envelopes can be computed using an algorithm that can sum the random variables corresponding to the return distributions of the portfolio segments. Such an algorithm should not assume that the random variables are independent (or have any other particular dependency relationship). The DEnv algorithm is one such algorithm (Zhang and Berleant 2005 [15]). That and others are described to varying degrees in Helton and Oberkampf (2004 [11]). We designate these envelopes "**Left**" and "**Right**" in Figure 6. The left and right envelopes enclose a space often called a p-box or probability box (e.g. Baudrit and Dubois 2005 [1]; the term is originally due to S. Ferson).

Continuing to develop our model, we have a particular concern with the left envelope, because it relates to worst case portfolio performance. Therefore we now associate $\alpha = 1$ with the envelope, henceforth ignoring the right envelope, which is also controlled by α and forms the other bound on the space of distributions representing our ignorance. Having defined the meanings for $\alpha = 0$ and $\alpha = 1$, we further detail the α parameterization by defining meanings for intermediate values of α. In doing this the concept of *horizontal averaging* is useful. Horizontal averaging takes two cumulative distributions $F_1(.)$ and $F_2(.)$, and returns a third distribution $F_{h\_ave}(.)$ which is midway between them in the sense that for each probability value $\text{Pr} \in [0, 1]$ on the *y*-axis,

$$F_{h\_ave}^{-1}(\text{Pr}) = \frac{F_1^{-1}(\text{Pr}) + F_2^{-1}(\text{Pr})}{2}$$ defines the corresponding *x*-coordinate of $F_{h\_ave}(.)$. An example

is the curve labeled "$\boldsymbol{\alpha = 0.5}$" in Figure 6. If there are places where $F_1^{-1}(\text{Pr})$ and $F_2^{-1}(\text{Pr})$ are undefined, i.e. there are horizontal segments in either $F_1(.)$ or $F_2(.)$, then the more general interval

extension can be used: $\mathbf{F}_{h\_ave}^{-1}(\text{Pr}) = \dfrac{\mathbf{F}_1^{-1}(\text{Pr}) + \mathbf{F}_2^{-1}(\text{Pr})}{2}$. We will henceforth use $F_1^{-1}(\text{Pr})$ and

$F_2^{-1}(\text{Pr})$ for expository simplicity, leaving interval extensions as an exercise for the interested reader. The horizontal averaging formula gives $F_1^{-1}(\text{Pr})$ and $F_2^{-1}(\text{Pr})$ equal weights of ½. Generalizing this to any pair of weights gives

$$F_{h\_ave}^{-1}(\text{Pr}) = \frac{\alpha F_1^{-1}(\text{Pr}) + (1-\alpha)F_2^{-1}(\text{Pr})}{2}. \qquad (6)$$

Figure 6 shows an example for $\alpha = 0.5$. Values of α near zero imply envelopes to the left of, but close to, the best-guess curve, representing little uncertainty about the shape of the curve. Values of $\alpha < 0$ are ruled out since there cannot be less than zero epistemic uncertainty. Values of α near one imply envelopes near, but to the right of, the left envelope. A value of $\alpha > 1$ in our development would imply ignorance of more than just lack of knowledge about the dependencies among segments. Certainly, other sources of ignorance do exist (although uncertainty about portfolio segment weights is not expressed with α but rather is handled with the mean $- z_* $ risk optimization process discussed earlier in this paper). If such a curve for $F_{h\_ave}(.)$ was in Figure 6, it would be to the left of the "**Left envelope, dependencies unknown**" curve.

Let r_b be a best-guess portfolio return distribution and L be a left envelope curve for it representing the case of $\alpha = 1$. Let L_α represent the horizontal average of L and r_b for a weight of α (Eq. 6). Then $r_b = L_0$ and $L = L_1$. Then we can seek the maximum value of α such that $L_\alpha \succ \widetilde{R}$ for some reference curve $\widetilde{R}$. Checking for FSD ($\succ_1$) involves verifying that curves L_α and $\widetilde{R}$ do not cross, while checking for SSD ($\succ_2$) requires integrating the distributions numerically and verifying that these integral curves do not cross. In either case the space of candidate best portfolios provided by the function $OPT(z)$ must be searched for a portfolio whose maximum α is at least as high as that of any other portfolio.

Example. The 3-segment portfolio example used throughout this paper was analyzed to find maximum values of α for different values of z under an SSD constraint. The results are shown in Table 4. Figure 7 shows the return distribution of the portfolio for $z = 3.98$ in more detail. The curve shown for $\alpha = 1.5894$ is the horizontal average (Eq. 6). It seems to cross the reference curve. In fact, it does cross it, thereby violating FSD. However if integrated, the resulting curves do not cross, so SSD holds, though just barely.

z	**max** α	μ
0.20	1.33	1.0800
1.00	1.33	1.0800
2.00	1.3380	1.0740
3.00	1.3410	1.0718
3.90	1.3408	1.0717
3.96	1.371	1.0706
3.97	**1.373**	1.0706
3.98	**1.373**	1.0706
3.99	1.371	1.0706
4.00	1.3700	1.0715
4.01	1.3680	1.0705
4.10	1.362	1.0705
5.00	1.36	1.0700

Table 4. The distribution for the return of the optimal portfolio given $z = 3.97$ or $z = 3.98$ has the most robust SSD over the reference curve, as measured by α (rather than |SSD| as in Section 2.3.2).

2.3.5 Search for the portfolio with the highest best-guess μ among portfolios having a left envelope with SSD over the reference curve (Table 2, item 4a). An investor may want a portfolio with as high an expected return as possible, but only if it meets some standard for

robustness to errors in its estimated return distribution. This objective seeks to combine two needs: robustness and high expected return. More formally, we seek a portfolio with mean return μ such that $\mu = \sup_{y \in Y, L_y \succ_2 \tilde{R}} \mu_y$, where $\tilde{R}$ is a reference curve and the supremum is over the portfolios whose left envelopes have SSD over $\tilde{R}$. Recall that L_y bounds the space of return distributions of portfolio y containing the distribution for each possible set of dependency relationships among the segments. Thus, requiring L_y to have SSD over $\tilde{R}$ ensures that return distribution r_y of portfolio y has SSD over $\tilde{R}$ regardless of the dependencies among segments in y. That is useful when we are not sure of what those dependencies are.
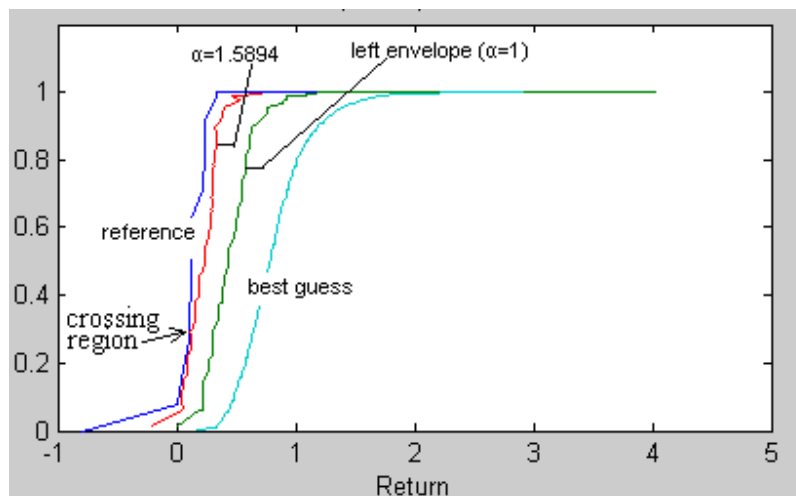

Figure 7. The curve for $\alpha = 1.5894$ models the presence of more epistemic uncertainty than the left envelope (i.e. the curve for $\alpha = 1$). It crosses the reference curve slightly, but the integrals of these curves do not cross, so SSD holds. The curves are for the optimal portfolio for $z = 3.98$.

In other words, we are searching for the best expected return available within the set of portfolios with return distributions whose left envelopes have SSD over the reference curve. This approach is like Table 2, item **2**, except in place of the "**Best-guess**" curve (Figure 6), it uses the "**Left envelope, dependencies unknown**" curve (Figure 6). This approach is also like Table 2, item **3**, except instead of maximizing α, it uses $\alpha = 1$ as a filter (recall that α is defined to equal 1 for the left envelope curve), and searches the portfolios that make it through the filter for one with the highest μ.

Example. In Table 5, optimal portfolios for each value of z all have left envelopes with SSD over the reference curve, as evidenced by $\alpha > 1$ in each case. Of these, the highest mean return is provided by the portfolio corresponding to z in the range [0, 1] (see Table 1), making that the best portfolio under this criterion.

15

Z	maximum α	μ
$z = 0.2$	1.3500	**1.0800**
$z = 1$	1.3500	**1.0800**
$z = 2$	1.5300	1.0740
$z = 3$	1.5894	1.0717
$z = 4$	1.5500	1.0705
$z = 5$	1.4000	1.0700

Table 5. Risk aversion coefficient (z) values, the robustnesses (maximum α) of the best portfolios corresponding to those values, and the mean returns of those portfolios.

2.3.6 Find a portfolio that maximizes μ, given a value of α to use as a filter (Table 2, item 4b). This criterion generalizes the one just described (item **4a** in Table 2) to enable seeking high return with an investor-determined degree of robustness. In the previous section we sought a portfolio that maximizes μ, and has a left envelope with SSD over the reference curve (i.e., $\alpha = 1$). In this section, we still seek a portfolio that maximizes μ, and wish to allow any value to be specified for α. An analyst who chooses a value for α is in effect stating an ignorance level (the epistemic uncertainty) corresponding to some subset (or superset) of the space of distributions associated with the range of possible dependency relationships among the portfolio segments. Complete ignorance about dependency implies $\alpha = 1$, and complete knowledge implies $\alpha = 0$. Intermediate degrees of ignorance imply intermediate values of α. Ignorance about the shape of the return distribution that includes additional factors besides the dependency relationships among the segments can be accounted for by increasing α, so $\alpha > 1$ is also allowed.

Let some value of α be given, representing the amount of epistemic uncertainty about the shape of a portfolio return distribution. Then any portfolio for which the horizontal average (by Eq. 6), with weight α, has SSD over the reference curve is eligible for consideration. Figure 6 illustrates such an eligible distribution for $\alpha = 0.5$. Any eligible distribution will have SSD over a given reference curve, with enough of a margin that even if its shape differs from the best-guess shape by as much as α permits, SSD still holds. From among the eligible portfolios, the one with the highest μ is considered best according to this criterion.

Example. Let the epistemic uncertainty associated with a best-guess portfolio return distribution be modeled as $\alpha = 1.5$. This implies the presence of other uncertainties besides uncertainty about the dependency relationships among the portfolio segment distributions. With such a high ignorance level, the horizontal average curve (Eq. 6) is even farther left than the left envelope. For the example we have been using (first presented in Section 2.2), some values of z result in portfolio return distributions whose horizontal averages have SSD over a reference curve chosen for this analysis, and some do not. Of those that do, the portfolio that is optimal for to $z = 2$ has a higher mean return μ than the others, and so is the best of those assessed in Table 6. Note that because of epistemic uncertainties in the shapes of the return distributions, the values for μ in Table 6 must be regarded as best guess values.

$\alpha = 1.50$		μ
z	|SSD|	
0.2	Negative	1.08
1	Negative	1.08
2	**Positive**	**1.074**
3	**Positive**	1.0717
4	**Positive**	1.0705
5	Negative	1.07

Table 6. A rather stringent robustness requirement expressed as $\alpha = 1.50$ filters out portfolios whose return distributions are too steep or too stretched out relative to a reference curve. Of the eligible portfolios tested ($z = 2$, 3 & 4), the highest mean return was 1.074.

2.3.7 Finding the expected value of information about α to help choose what α to use (Table 2, item 4c). Recall the strategy in the lower half of Table 2 of maximizing mean return μ involves, first, identifying portfolios that meet an SSD requirement, and second, searching those portfolios for one with the highest possible μ. In that light, let us consider next how obtaining more information about the shape of a portfolio return distribution r (i.e. decreasing epistemic uncertainty about it) will tend to increase the chance that it will have SSD over a reference curve (i.e. that $r \succ_2 \tilde{R}$). Such information will tend to move left envelopes rightward, consequently making more portfolios eligible. That in turn will tend to raise the maximum μ available due to the enlarged set of eligible portfolios.

Figure 8 illustrates a "**Left envelope, dependencies unknown**" curve without FSD over reference curve $\tilde{R}$. However, its "**Left envelope, α = 0.5**" curve does have FSD over $\tilde{R}$. Thus, reducing epistemic uncertainty about the "**Best-guess**" curve from $\alpha = 1$ to $\alpha = 0.5$ moved the portfolio represented by the "**Best-guess**" curve into the set of qualified portfolios. If this portfolio happened to have the highest available μ, the value v of the information that reduced α to 0.5 would be $v = \mu_{newmax} - \mu_{oldmax}$, where μ_{newmax} is the expected return of this newly qualified portfolio and μ_{oldmax} is the expected return of the best of the qualified portfolios prior to reducing α to 0.5. The generalization to SSD is straightforward; the situation for FSD is shown because it visualizes well.

We can determine the "demand value" of information about α (Ben-Haim 2006 [2]) from a plot of α vs. μ. Figure 9 shows a schematic example. L_y, the left envelope for portfolio y assuming $\alpha = 1$, has parameterized form $L_{y, \alpha=k}$. The more uncertainty in the shape of portfolio return distributions (expressed as larger values of α), the more leftward are their left envelopes, hence the fewer portfolios have left envelopes with SSD over $\tilde{R}$. This tends to lower the maximum expected return available from among them.

Function $f(.)$ in the figure can be used to determine the incremental value of obtaining information that reduces α from α_2 to α_1. That value is $\Delta v = f(\alpha_1) - f(\alpha_2)$ where α_2 is the current value of α and α_1 is a smaller (i.e. more informative and thus useful) value. If this information costs below Δv, the expenditure is worth making. An important special case is reducing α from 1 to 0, that is, going from no information about the dependencies among portfolio segments to fully defining the

dependencies among the portfolio segments. The cost of this information is worth paying if it is below $f(0) - f(1)$.
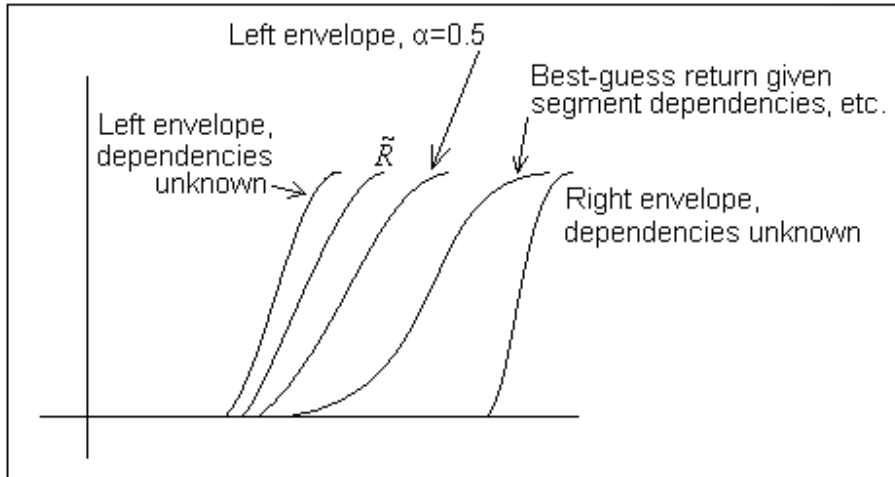

Figure 8. The best-guess distribution has FSD over $\tilde{R}$. If $\alpha = 1$ then the "**Left envelope, dependencies unknown**" curve applies and FSD might not hold, depending on what the true but unknown distribution actually is. On the other hand, if $\alpha = 0.5$, FSD does hold.
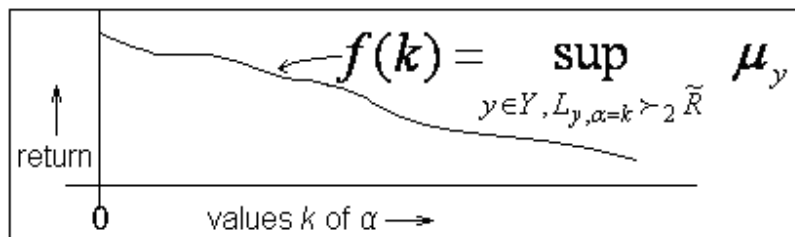

Figure 9. Maximum expected return μ over all portfolios y whose left envelopes L_y stochastically dominate reference curve $\tilde{R}$.

3 Conclusion

This paper introduces an approach, and specific variations (Table 2), to determining the best possible investment plan given the two standard conflicting portfolio investment goals of mean return and risk. On the one hand we seek a high expected (mean) return. On the other we seek to control risk. To manage risk we seek to guarantee that the portfolio model has second-order stochastic dominance (SSD) over a minimum tolerable reference curve, because it has been shown that (1) if an SSD relationship exists between two return distributions, any risk-averse investor will prefer the dominant one, and (2) this constraint is weaker than the FSD relationship, which is unnecessarily strong. Strong constraints are undesirable because they reduce the space of allowable portfolios, tending to limit investment choices.

We find the best portfolio by first generating a set of optimal portfolios. Then we search the set using stochastic dominance and Information-Gap Theory to identify the best one. The traditional approach to portfolio optimization using Markowitz theory is challenged when correlations or other dependencies among portfolio segments are hard to provide, return distribution shapes are uncertain, there is a lack of price data, or various other fundamental data are unavailable. The analyses shown in this paper address the first two of these challenges, thereby showing how rational portfolio choice is possible even under severe (epistemic) uncertainty.

4 Acknowledgements

5 References

[1] C. Baudrit and D. Dubois, Comparing methods for joint objective and subjective uncertainty propagation with an example in a risk assessment, *Proceedings of the Fourth International Symposium On Imprecise Probabilities and Their Applications* (*ISIPTA '05*), Carnegie Mellon University, Pittsburgh, July 20-23 2005, pp. 416-422, www.sipta.org/isipta05/proceedings.

[2] Y. Ben-Haim, *Information-Gap Decision Theory*, 2nd ed., Academic Press, 2006.

[3] D. Berleant, M. Dancre, J.-P. Argaud and G. Sheble, Electric company portfolio optimization under interval stochastic dominance constraints, *Proceedings of the Fourth International Symposium On Imprecise Probabilities and Their Applications* (*ISIPTA '05*), Carnegie Mellon University, Pittsburgh, July 20-23 2005, pp. 51-57. www.sipta.org/isipta05/proceedings

[4] Z. Bodie, A. Kane and A.J. Marcus, *Investments*, 4th ed., Irwin/McGraw-Hill, Chicago, 1999.

[5] F.D. Chabi-Yo, Stochastic discount factor volatility, bound and portfolio selection under higher moments, *44e Congrés annuel de la Société canadienne de science économique*, May 5-6, 2004, Quebec, www.scse.ca/scse/congres2004.

[6] M.-P. Cheong, D. Berleant and G. Sheblé, Information-Gap Decision Theory as a tool for strategic bidding in competitive electricity markets, *Proceedings of the 8th Int. Conf. on Probabilistic Methods Applied to Power Systems*, Sept. 12-16, 2004, Ames, Iowa.

[7] E. De Giorgi, Reward-risk portfolio selection and stochastic dominance, *Journal of Banking and Finance*, **29** (2005), 895-926.

[8] E.J. Elton, M.J. Gruber, S.J. Brown and W.N. Goetzmann, *Modern Portfolio Theory and Investment Analysis*, 6th ed., John Wiley & Sons, New York, 2003.

[9] S. Ferson, *Ramas Risk Calc: Risk Assessment with Uncertain Numbers*, CRC Press LLC, 2002.

[10] S.J. Grossman and R.J. Shiller, The determinants of the variability of stock market prices, *The American Economic Review*, **71** (1981), 222-227.

[11] J.C. Helton and W.L. Oberkampf, eds., Special Volume on Alternative Representations of Epistemic Uncertainty, *Reliability Engineering and System Safety*, **85** (1-3) (July-Sept. 2004), 1-369.

[12] H. Markowitz, Portfolio selection, *Journal of Finance,* **7** (1) (1952), 77-91. Reprinted as Cowles Commission Report, New Series, No. 60, www.ifa.com/Media/Images/PDF%20files/PortfolioSelectionMarkowitz.pdf.

[13] P. Perny, O. Spanjaard, and L.-X. Storme, State space search for risk-averse agents, *Int. Joint Conf. on Artificial Intelligence* (*IJCAI* '07), Hyderabad, Jan. 6-12, pp. 2353-2358. www.ijcai.org/papers07/Papers/IJCAI07-379.pdf

[14] R.T. Rockafeller, *Convex Analysis*, Princeton U. Press, 1970.

[15] J. Zhang and D. Berleant, Arithmetic on random variables: squeezing the envelopes with new joint distribution constraints, *Proceedings of the Fourth International Symposium On Imprecise Probabilities and Their Applications* (*ISIPTA '05),* Carnegie Mellon University, Pittsburgh, July 20-23 2005, pp. 416-422.